



Europäisches Patentamt  
European Patent Office  
Office européen des brevets



(11) Publication number : **0 634 854 A2**

(12)

## EUROPEAN PATENT APPLICATION

(21) Application number : **94304973.4**

(51) Int. Cl.<sup>6</sup> : **H04L 12/56**

(22) Date of filing : **06.07.94**

(30) Priority : **13.07.93 US 90732**

(43) Date of publication of application :  
**18.01.95 Bulletin 95/03**

(84) Designated Contracting States :  
**DE FR GB**


(71) Applicant : **AT & T Corp.**  
**32 Avenue of the Americas**  
**New York, NY 10013-2412 (US)**

(72) Inventor : **Karol, Mark John**  
**26 Poplar Avenue**  
**Fair Haven, New Jersey 07704 (US)**

(74) Representative : **Buckley, Christopher Simon**  
**Thirsk et al**  
**AT&T (UK) LTD.,**  
**AT&T Intellectual Property Division,**  
**5 Mornington Road**  
**Woodford Green, Essex IG8 0TU (GB)**

(54) **Controlled-feedback packet switching system.**

(57) Packets or cells received from different input ports ( $101_1, 101_n$ ) of a switch (102) and destined for a common output port ( $103_1 - 103_n$ ) of that switch, are analyzed to determine their priority level. Lower-priority packets or cells are buffered in recirculation delay lines ( $105_1 - 105_m$ ) of appropriately-selected lengths, and thereafter scheduled for transmission to the output port based on their level of priority.



EP 0 634 854 A2

## Technical Field

This invention relates to switching systems. More specifically, this invention relates to the scheduling of packet delivery in a switching system.

## Background

Because of the unscheduled nature of arrivals of packets or ATM cells to a packet switching system, two or more packets may simultaneously arrive on different inputs destined for the same output. The switch architecture may allow one of these packets to pass through to the output, but the others must be queued for later transmissions. This temporary congestion caused by simultaneous arrival of packets or cells is typically handled by temporarily storing the packets or cells in buffers. For traffic distributions that are random or more or less uniform, buffering requirements are rather lenient. However, for high performance packet switching systems designed to handle bursty traffic, the buffering requirements are more stringent.

For electronic packet switches, buffering is ordinarily implemented in a random access memory (RAM) that is typically shared by all the inputs and outputs of the switch in order to reduce memory storage requirements. In optical packet switches, the present lack of an optical random access memory significantly complicates buffering in those optical switches. Approaches that have been considered for buffering in optical switches include an implementation that involves the routing of queued packets to trap lines that retard the transmission of the queued input packets to the desired output, thereby allowing other input packets destined for the same output to be transmitted during the delay period. However, this approach presents certain drawbacks that prevent its use in optical and optoelectronic packet switches. Specifically, certain scheduling functions needed for the orderly and timely switching and transmission of packets are not performed in the trap line approach. For example, the trap line approach (unlike the RAM approach) does not permit changes to the "scheduled" transmission time of lower-priority packets when higher-priority packets arrive later.

Another approach that has been advocated for buffering packets in optical switches is the so-called "feed-forward" technique in which packets contending for an output port are delayed by different numbers of time slots to avoid collision with previously scheduled packets. In that approach, packets are dropped if they cannot be scheduled in a collision-free manner. This approach, however, does not allow transmission time to be updated on a slot-by-slot basis and does not adequately support priority traffic.

Thus, there is a need for a packet buffering system for use in optical and optoelectronic packet

switches which offers the same performance and functionality provided by a RAM in electronic packet switches.

## Summary

This invention is directed to a packet switch in which incoming packets or cells destined for a common output port are analyzed to determine their priority level. Lower-priority packets or cells are buffered in recirculation delay lines of appropriately-selected lengths, and thereafter are scheduled for transmission to the output ports based on their level of priority.

In a specific example of the invention, a memory-less non-blocking switch is designed to include an input/output section comprised of a certain number of input ports and output ports and a certain number of recirculation delay lines of various lengths connected to dedicated input and output ports for buffering packets. The switch also includes control circuitry that determines which packets need to be buffered in the recirculation delay lines and schedules the delivery of the buffered packets based on their priority level. The control circuitry also keeps packets in their proper first-in, first-out sequence, supports multiple levels of priority traffic, and ensures that packets pass through the recirculation delay lines only a small number of times to minimize power losses, thereby avoiding the need for optical amplifiers in the delay lines in most cases. If amplifiers are used, a reduced number of lines through which a packet has to travel results in a proportional decrease in amplifier noise.

In another example of the invention, a packet switching system is partitioned into multiple, memory-less, non-blocking switches that are either connected to recirculation delay lines or to the output ports of the packet switching system. The first group of switches are called a "delay-line switches" while the second are called an "output switches". Delay-line and output switches are front-ended by routers which receive commands from a control circuitry to direct packets from input ports to either a delay-line switch when the packets must be queued or to an output switch when no queuing is required.

## Brief Description of the Drawing

FIG. 1 illustrates a packet switch designed in accordance with the invention to schedule delivery of lower-priority packets based on their priority level and to buffer packets in recirculation delay lines of appropriately-selected lengths;

FIG. 2 shows one example of a control table which keeps track of packets in the recirculation delay lines of FIGs. 1 and 3; and

FIG. 3 shows an illustrative packet switching system that is partitioned into multiple switches

front-ended by routers.

### Detailed Description

FIG. 1 illustrates an example of a packet switch which schedules delivery of packets based on their priority level. The packet switch of FIG. 1 also buffers lower-priority packets in recirculation delay lines of appropriately-selected lengths. FIG. 1 shows a block diagram of a  $n \times n$  packet switch comprising a)  $n$  input lines 101-1 to 101- $n$  b)  $n$  output ports 103-1 to 103- $n$  c)  $m$  recirculation delay lines (of appropriately-selected lengths) 105-1 to 105- $m$  for buffering packets d) an  $(n+m) \times (n+m)$  memoryless, non-blocking switching block 102, and e) a control circuit 104 which reconfigures the switching block 102 on a packet-by-packet basis. In FIG. 1, packets arrive at the input lines 101-1 to 101- $n$  of the switching block 102. While switching block 102 can be a strictly non-blocking switching block, it is preferably a rearrangeably non-blocking switching block. For optical implementations, input lines 101-1 to 101- $n$  may be fiber lines. When there is no contention for the output ports, packets received from input lines 101-1 to 101- $n$  are transmitted to the appropriate output ports 102-1 to 102- $n$  based on the packet headers that indicate the destination point for each packet. The header of a packet is read by control 104 which uses power splitters 110-1 to 110- $n$  to tap a small fraction of the packet energy to read the header. To determine the appropriate configuration of the switching block 102, control 104 keeps track of all packets buffered in the recirculation delay lines 105-1 to 105- $m$ , so that it knows which packets will be returning to the switching block 102 at what times. Control 104 may be implemented using, for example, a microprocessor which executes programming instructions described below.

FIG. 1 shows that the recirculation delay lines 105-1 to 105- $m$  have lengths  $d_1, d_2, \dots$ , and  $d_m$ . The lengths of those lines are expressed in units equal to the number of packets that they respectively can store end-to-end. In order to explain the interaction between switching block 102, control 104 and recirculation lines 105-1 to 105- $m$ , it is assumed in this example, (without loss of generality, and for the sake of simplicity) that  $d_1 \leq d_2 \leq \dots \leq d_m$ . For illustrative purposes, special attention is devoted to the case in which  $d_1 = 1, d_2 = 2, \dots$ , and  $d_m = m$ . Since a total of  $B \triangleq (d_1 + d_2 + \dots + d_m)$  packets can be stored in the recirculation delay lines 105-1 to 105- $m$ , it follows that  $B = m(m+1)/2$  when  $d_1 = 1, d_2 = 2, \dots$ , and  $d_m = m$ . It is further assumed that switching block 102 and control circuit 104 use a time-slotted system in which, for each time slot, up to  $(n+m)$  packets may arrive at the switch (i.e.,  $n$  new arrivals plus  $m$  "recirculation packets"). Of these packets, up to  $n$  can be transmitted to their appropriate output ports; the rest are "buffered"

in the recirculation delay lines 105-1 to 105- $m$ . A packet that is buffered in a recirculation delay line of length  $d_i$  will exit that delay line and return to switching block 102 after  $d_i$  time slots. Control 104 selects which packets to output for each time slot, and also assigns the remaining packets to the recirculation delay lines. More specifically, control 104 assigns packets to the appropriate delay lines in a way such that "buffer locations" in the delay lines are efficiently utilized, and packets retain their proper first-in, first-out sequence, as needed. Control 104 also ensures that each packet circulates through the delay lines 105-1 to 105- $m$  only a small number of times. This last property is very important, because it may help to keep the power losses low enough to reduce, and perhaps eliminate, the need for optical amplifiers in the recirculation delay lines 105-1 to 105- $m$ . If amplifiers are needed in the recirculation delay lines 105-1 to 105- $m$ , minimization of the number of recirculations helps keep the added noise sufficiently small. In addition, if the power budget does indicate the need for amplifiers in the recirculation delay lines 105-1 to 105- $m$ , the reduction in the number of lines translates into a valuable reduction in the number of amplifiers.

It may be important to note that many different technologies can potentially be used for an optical implementation of the switch illustrated in FIG. 1. For example, technologies such as guided-wave using lithium niobate; or perhaps Wavelength Diversion Multiplexing (WDM) and star couplers; or even wavelength routers can be used to build such a switch. In addition, the delay-line function that is performed in this example by the recirculation delay lines 105-1 to 105- $m$  can be implemented using, for example, optical fiber lines or other appropriate optical transmission means. As to the control circuitry 104, it can more easily be implemented using a microprocessor even though it is also possible to use optical logic gates to design such a circuit in an optical implementation.

FIG. 2 shows one exemplary structure of a control table which is arranged to keep track of packets in the recirculation delay lines of FIG. 1 described above and FIG. 3 described below.

In order to free control circuit 104 of FIG. 1 of the task of reading (for each time slot) the headers of all packets exiting the recirculation delay lines 105-1 to 105- $m$ , control circuitry 104 maintains a table of information about the packets stored in the recirculation delay lines 105-1 to 105- $m$ . FIG. 2 illustrates one possible way for control 104 to store information about the buffered packets. As mentioned above, this information (e.g., the packets' output port addresses) is obtained by reading the packet headers when they first arrive to the switching block 102. This information is kept in the control table of FIG. 2 until the packet exits the switching block 102. In this example, the table of FIG. 2 is an  $m \times d_m$  table in an (electronic) RAM,

that operates at the packet rate and that is arranged to mimic the flow of packets through the recirculation delay lines 105-1 to 105-m. Each row in the table corresponds to one of the m delay lines ( $m = 8$  in FIG. 2). Each column contains information corresponding to the set of packets that will exit the delay lines at the same time and arrive together for routing through the switching block 102. The time-slot pointer shown at the bottom of FIG. 2 cyclically shifts to the left one column per time slot to indicate the current set of (up to) m recirculation packets now returning to switching block 102. The shaded portion of the  $m \times d_m$  table corresponds to unused entries (for a given position of the time-slot pointer). When a packet is buffered in a recirculation delay line, information about that packet is moved in the control table of FIG. 2 to the appropriate position along the "diagonal" ( $P_{14}$  for delay line 1,  $P_{25}$  for delay line 2,  $P_{36}$  for delay line 3, ..., or  $P_{83}$  for delay line 8). The amount of information associated with each of these  $P_{ij}$  entries depends on the particular implementation of control 104, as illustrated by the following two examples.

In an illustrative non-FIFO (First-In First-Out) control implementation, the control 104 first routes as many of the packets to their outputs as possible for each time slot. In sequence, it considers the recirculation packets (beginning at the longest delay line and moving towards the shortest delay line), and then considers the new packet arrivals. If there are multiple priorities, higher-priority packets are handled first. Whenever a packet cannot be routed to its output port j (because another packet has already been selected for routing to output j), then the packet is sent to the shortest delay line that has the fewest packets destined for that output j "in that column" of the Control Table of FIG. 2. Since at most one packet per output can be transmitted for each time slot, this last factor helps "load balance" the output addresses over the table's columns.

In an illustrative FIFO-control implementation, for each time slot, the new packet arrivals and the recirculation packets are either (i) routed to the appropriate output ports for transmission, (ii) "scheduled" for transmission after one more recirculation, or (iii) left "unscheduled" and sent to a delay line for another recirculation. Each  $P_{ij}$  entry in the control table keeps track of (i) the input-output ports of each packet in the recirculation delay lines, and (ii) whether or not the packet is scheduled for transmission the next time it reaches the switching block 102. There also is a FIFO table associated with each input-output pair. The FIFO keeps a first-in, first-out list of all packets of this input-output pair, the exact location of these packets in the recirculation delay lines and in which (if any) future time slots they are scheduled for transmission. Finally, a "timestamp" is given to each packet when it first arrives to the switching block 102. The timestamp corresponds to the transmission time of an ideal out-

put-queuing switch, and represents the packet's "anticipated transmit time."

Using this information about each packet, the Control schedules packets for transmission in the current time slot or future time slots. First, it routes any "scheduled packets" on the recirculation delay lines to the appropriate output ports. Second, it checks if any of the new arrivals can be routed to their outputs (without violating the FIFO constraint). Third, it schedules packets, if possible, for transmission after their next recirculation (on the shortest possible delay line). A packet can be scheduled only if the packet before it in its input-output FIFO has already been scheduled. Priority in this scheduling is given to packets with the smallest timestamp ("anticipated transmit time"), and to recirculation packets on the longest delay lines. Once a packet is scheduled, this may also indirectly allow other packets waiting in the same input-output FIFO to be scheduled without violating the FIFO constraint. Finally, any remaining packets are left unscheduled and are recirculated to the delay lines. Starting with unscheduled packets having the smallest timestamp, packets are sent to the shortest delay line that has the fewest packets destined for that output in that column of the control table of FIG. 2.

Other contention resolution schemes considered include routing a packet to an output port (as opposed to a recirculation delay line) based on the position of the input port from which the packet is received when that packet is contending for an output port with another packet of equal priority level. Likewise, a packet received from a recirculation delay line dedicated input port may be routed to an output port when another packet received simultaneously from a "regular" input port is contending for the same output port. In that case, in accordance with the principles of the invention, the other packet would be routed to a recirculation delay line. priority

FIG. 3 shows an illustrative packet switching system that is partitioned into multiple switches frontended by routers. The modular design, illustrated in FIG. 3, partitions a packet switch into a "memory block" comprised of memoryless non-blocking switching blocks (MNSB) 306, 312, 314, and 330, delay lines 318-1 to 318-m, and an input/output section comprised of input lines 301-1 to 301-n and output ports 340-1 to 340-n, respectively. In FIG. 3, the  $(n+m) \times (n+m)$  switch of FIG. 1 is configured as  $n \times n$ ,  $n \times m$ ,  $m \times n$ , and  $m \times m$  MNSBs, plus a number of routers 303, 304, 319, 320 and optional rear-end switches 309, 310, 316 and 317 for certain implementations.

The packet switching system that is illustrated in the block diagram of FIG. 3 includes routers 303, 304, 320 and 319 that directs packets to either a delay-line MNSB, such as MNSB 314 (330) or an output MNSB such as MNSB 312 (306). Each MNSB is connected to n routers that receive incoming packets from input

lines 301-1 to 303-n or m routers that receive recirculated packets from recirculation delay lines 318-1 to 318-m. In an optical implementation of the packet switching system of FIG. 3, routers 303 and 304 also perform a power splitting function that allows a control 307 to tap a small fraction of a packet's energy to read that packet's header. The packet switching system that is illustrated in the block diagram of FIG. 3 shares some common elements with the packet switch shown in FIG. 1. For example, the packet switching system of FIG. 3 uses a) the same set of recirculation delay lines shown in FIG. 1, and b) a similar control structure, albeit with a slightly more distributed implementation. Thus, the control table illustrated in FIG. 2 is readily applicable to the modular design of FIG. 3.

As is the case for the switch of FIG. 1, the Control 307 is responsible for all routing decisions, including deciding which delay lines buffer which packets. For example, when multiple incoming packets destined for the same output port are received by routers 303 or 304 from input lines 301-1 to 301-n, the headers of these packets are read by control 307 which decides the proper treatment for the received packets using, for example, the scheduling techniques described above. As part of that decision, control 307 determines which packet among the received packets to send to the appropriate output port based on the priority level of the packets. The selected packet to be routed to the output port is sent to MNSB 306 which promptly forwards that packet to rear-end switch 309 or 310 via line 308. Rear-end switch is used in optical implementations of the packet switching system of FIG. 3 to prevent energy loss that could result from the use of other types of devices. When a packet has to be buffered, control 307 sends a signal to router 303 or 304 to forward that packet to MNSB 330 which in turn sends the packet to the delay line indicated by control 307.

Advantageously, buffering packets in the recirculation delay lines (105-1 to 105-m in FIG. 1 or 318-1 to 318-m in FIG. 3) offers flexibility and efficiency not possible with other techniques. For example, if first-in first-out requirements dictate that an arriving packet cannot be transmitted for at least k time slots (because other packets are queued for the same output), then, it is advantageous to store that packet in a recirculation delay line of length k (if possible) and transmit that packet in due time. Even if the packet cannot be buffered in a delay line of length k, "optimal performance" may still be possible since there may be many combinations of delay-line lengths that sum to k. For example, an optimal (i.e., minimal) delay of 10 time slots can be attained by buffering a packet in a delay line of length 10, or by buffering it successively in delay lines of length 7 and 3, or successive delay lines of length 6, 1, and 3, etc. Furthermore, recirculation permits "scheduling decisions" to be revised

each time a packet returns to the switch from the feedback delay lines (e.g., to allow for quicker transmission of new higher-priority packets).

## Claims

1. A packet switching system comprising:
  - at least one switching block having a plurality of input and output ports;
  - a control circuit for determining a) priority levels of packets received by the input ports and b) delivery time of packets to the output ports;
  - at least one recirculation delay element;
 and
  - means for buffering packets in the at least one recirculation delay element in accordance with the priority levels determined by the control circuit.
2. The invention of claim 1 wherein said buffering means comprise
  - means for routing at least one lower-priority packet to the at least one recirculation delay element until the lower-priority packet is no longer in contention for the output port with at least one higher-priority packet; and
  - means for keeping track of all packets routed to all recirculation delay elements.
3. The invention of claim 2 further comprising
  - means for limiting the number of times a lower-priority packet is routed to a recirculation delay element.
4. The invention of claim 1 further comprising
  - means for routing a lower-priority packet to the output port after at least one higher-priority packet has been routed to said output port.
5. The invention of claim 1 wherein the packet switching system has more than one recirculation delay elements of different time-delaying means and wherein a packet is buffered in a particular recirculation delay element based on its priority level.
6. The invention of claim 5 wherein lower-priority packets are buffered in longer-time-delaying-means element and higher-priority packets are buffered in shorter-time-delaying-means element.
7. The invention of claim 1 wherein each recirculation delay element is a fiber-optic line that is of a particular length and that is connected to a dedicated input and output port of a switching block.

8. The invention of claim 1 further comprising  
means for keeping all packets routed to an output port of the switching block in a first-in first-out sequence.
9. Apparatus for use in a packet switching system comprising:  
means for determining the level of priority of incoming packets received from different input ports of the packet switching system and destined for a common output port of the packet switching system;  
routing a higher-priority packet to the output port;  
means for buffering in at least one recirculation delay line at least one lower-priority packet; and  
means for scheduling delivery to the output port of the at least one lower-priority packet buffered in the recirculation delay lines.
10. The apparatus of claim 9 wherein the scheduling means comprises:  
means for routing at least once the at least one lower-priority packet to the recirculation delay lines.
11. The invention of claim 10 further comprising  
means for limiting the number of times a lower-priority packet can be routed to the at least one recirculation delay line.
12. A method of routing packets in a switch comprising  
determining levels of priority of incoming packets received from different input ports of the switch and destined for a common output port of the switch;  
routing a higher-priority packet to the output port;  
buffering in at least one recirculation delay line at least one lower-priority packet; and  
scheduling delivery to the output port of the at least one lower-priority packet buffered in the at least one recirculation delay line.
13. A packet switching system comprising:  
a switching fabric having at least one switching block with a plurality of input and output ports;  
a control circuit for determining levels of priority for packets received from the input ports;  
at least one recirculation delay element which is connected to a dedicated input port and a dedicated output port of the switching fabric and through which, a particular packet exiting the switching fabric is returned to the switching fabric at least once; and
- means for queuing at least one packet in the at least one recirculation delay element in response to a signal from the control circuit.
14. The invention of claim 13 wherein the signal is emitted by the control circuit if at least two packets are received simultaneously on different input ports and are destined for a common output port.
15. The invention of claim 13 wherein the signal from the control circuit is a function of the priority levels of the received packets and programming instructions in said circuit.
16. The invention of claim 15 further comprising  
means for routing at least one lower-priority packet to the at least one recirculation delay element at least once.
17. The invention of claim 15 further comprising  
means for routing a queued packet to an output port when the queued packet is in contention for the output port with another packet of equal priority level.
18. The invention of claim 15 further comprising  
means for routing a packet to an output port based on the position of the input port from which said packet is received when said packet is in contention with another packet of equal priority level.






FIG.1

FIG. 2





*FIG. 3*



**THIS PAGE BLANK (USPTO)**

---

**This Page is Inserted by IFW Indexing and Scanning  
Operations and is not part of the Official Record**

**BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☒ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: \_\_\_\_\_

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.**

THIS PAGE BLANK (USPTO)